arXiv:1402.0928v3 [cs.DL] 5 Oct 2014

A Framework for Evaluation of Composite Memento Temporal Coherence

Scott G. Ainsworth
Old Dominion University
Norfolk, VA, USA
sainswor@cs.odu.edu

Michael L. Nelson
Old Dominion University
Norfolk, VA, USA
mln@cs.odu.edu

Herbert Van de Sompel
Los Alamos National Laboratory
Los Alamos, NM, USA
herbertv@lanl.gov

ABSTRACT

Most archived HTML pages embed other web resources, such as images and stylesheets. Playback of the archived web pages typically provides only the capture date (or Memento-Datetime) of the root resource and not the Memento-Datetime of the embedded resources. In the course of our research, we have discovered that the Memento-Datetime of embedded resources can be up to several years in the future or past, relative to the Memento-Datetime of the embedding root resource. We introduce a framework for assessing temporal coherence between a root resource and its embedded resource depending on Memento-Datetime, Last-Modified datetime, and entity body.

1. INTRODUCTION

Web archives, such as the Internet Archive [9], make best effort attempts to archive all or parts of the World Wide Web. Although these archiving efforts are easily viewed as the virtual analog of a traditional library that collects books and periodicals, the differences between archiving the products of printing presses and archiving the pages of the Web are many. This report focuses on the web archiving impact of a critical difference between composite physical and web products: when composition occurs. When a book or magazine is produced, the text, images, and other resources are collected and composed into a product by the publisher. Although future editions may use updated versions of the resources, a particular copy archived by a library is forever unchanging. Future library patrons, when viewing the copy, need not consider if the copy currently in their hands changed after it was produced. Thus, archiving physical media is deterministic from both the library's and patron's perspective.

In contrast to a book, web pages are composed much later in the production process: when it is viewed. Returning to a library book, it is as if the book's photographs were not included until after the book is opened. Web archives generally capture web resources using web crawlers such as Heritrix [8]. Due to resource constraints (e.g., network bandwidth) and "politeness" considerations (not unduly burdening original resource servers), archive crawls seldom capture a web page and its component resources simultaneously. Spaniol et al. [11] note that crawls may span hours or days, increasing the risk of temporal incoherence because resources may change during the crawl. Therefore, archiving composite web resources in a coherent state is probabilistic instead of deterministic.

Still, when an archived composite resource is presented in a web browser, it is labeled with the singular datetime of the root resource, as circled in figure 1. Most users would consider the presentation

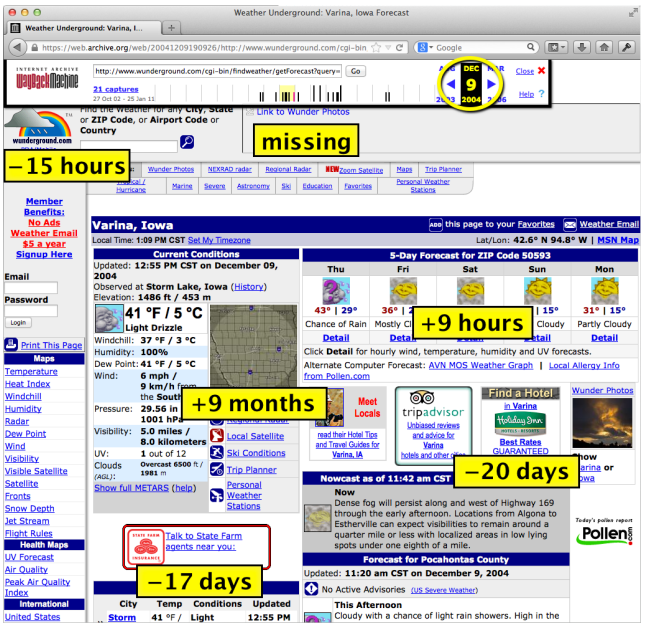

Figure 1: Embedded resource capture deltas, but entire page marked with a single datetime

a coherent representation of the composite resource as it existed at that displayed datetime. However, the capture datetime of the embedded resources can vary greatly. Figure 1 also shows the capture deltas for five embedded resources for an archived wunderground.org forecast page that was captured by the Internet Archive on 2004-12-09 19:09:26Z. The yellow rectangles indicate the capture deltas. Negative means captured before the root; positive, after the root. Note the *clear weather* satellite image that was captured +9 months after the root. Contrast it with the *chance of rain* and *mostly cloudy* images captured just 10 hours after the root. This disagreement reveals a temporal incoherence—specifically, the satellite image is *prima facie violative*.

In our research data, it is common for the temporal spread between the oldest and newest captures to be weeks or months. Unexpected though was the discovery of many composite resources with spreads exceeding one year. Even more unexpected were spreads exceeding five years; a few even exceed ten years.

Spaniol et al. [12], Denev [5], and Ben Saad et al. [3, 2] all introduce strategies to improve the quality of future web crawler-based

captures. If implemented, these strategies should reduce spread in new captures. In contrast, our work focuses on the quality of existing holdings, without the benefit of improved crawl strategies.

Our investigation into the causes and consequences of temporal spread has revealed two important things about embedded resources:

- Even if captured within seconds of the root resources, embedded resources are not always temporally coherent.
- Even if captured much later than the root resource, embedded resources are not necessarily incoherent.

Evaluation of temporal coherence depends on capture datetime, last modified datetime, and content differences. Together, these attributes form patterns which can be considered prima facie coherent, prima facie violative, possibly coherent, and probably violative. In order to facilitate reasoning about temporal coherence, we catalog embedded resource coherence patterns and the conditions that differentiate them.

The terminology used in this report comes primarily from the Memento framework [14, 15, 13], which enables time-based HTTP access to archived resources. The following definitions are from RFC 7089 [13]:

Original Resource An Original Resource is a resource that exists or used to exist, and for which access to one of its prior states may be required.

Memento A Memento for an Original Resource is a resource that encapsulates a prior state of the Original Resource. A Memento for an Original Resource as it existed at time T is a resource that encapsulates the state the Original Resource had at time T.

Memento-Datetime The "Memento-Datetime" response header is used by a server to indicate that a response reflects a prior state of an Original Resource. Its value expresses the datetime of that state.

TimeMap A TimeMap for an Original Resource is a resource from which a list of URIs of Mementos of the Original Resource is available.

URI-R The URI of an Original Resource.

URI-M The URI of a Memento.

URI-T The URI of a TimeMap.

At first glance, the Memento-Datetime header appears to duplicate the Last-Modified header; this is not the case. Last-Modified is set by an original resource's server and indicates the last datetime the resource changed. Memento-Datetime is the datetime the original resource was captured and archived. The existence of the Memento-Datetime header also entails a promise that the original resource's state is archived, frozen in time [13]. See the appendix for additional details and example headers.

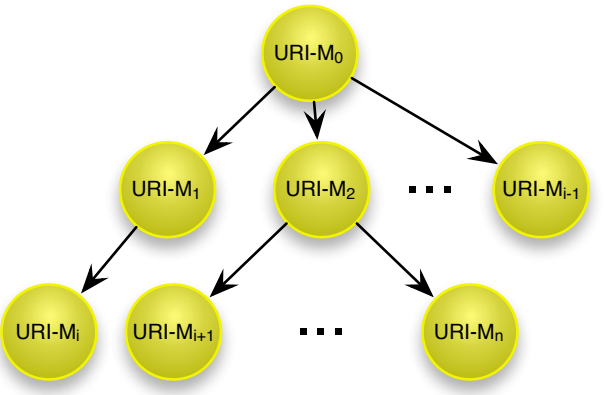

Figure 2: Composite Memento

Table 1: Definitions

Term	Definition
R	An original resource URI. A URI-R. R_0 is the root URI-R. R_i, $i \geq 1$, is an embedded URI-R.
M	A URI for an archived copy of a URI-R. A memento URI. A URI-M. M_0 is the root URI-M. $M_{i,j}$, $1 \leq j \leq n$, is a URI-M for R_i.
m	The representation of M. m_0 is the representation of M_0. $m_{i,j}$ is the representation of $M_{i,j}$.
$\mathbb{T}$	A timemap for R. A list of mementos for R. $\mathbb{T}_i = \{M_{i,j} \vert 1 \leq j \leq n\}$; n is the number of Ms archived for R_i.
$\mathbb{C}$	A composite memento. A set of M comprising the root URI-M, M_0, and zero or more embedded URI-Ms, $M_{i,j}$.
t	A target datetime (*Accept-Datetime* in RFC 7089 [13]).

2. COMPOSITE MEMENTOS

2.1 Definition

A *composite memento* is a root URI-M and all embedded URI-Ms required to recompose the presentation at the client. A composite memento generally comprises a root HTML resource and embedded resources such as images and stylesheets. Some embedded resources (e.g., HTML frames, style sheets) can also have embedded resources, which are also part of the composite memento. Figure 2 is a tree representation of a composite memento. (Technically, a composite memento is a graph, but in this context can treated as a tree without loss of generality.)

2.2 Recomposition

Recomposing a composite memento is the recursive process of selecting URI-Ms for URI-Rs, retrieving the representations for the URI-Ms, and extracting URI-Rs embedded in those URI-M representations. The variables in table 1 and functions in table 2 are used in the description that follows. R is simply a short form for URI-R, with a root URI-R designated R_0 and embedded URI-Rs designated R_i, $i > 0$. Likewise, M is short for URI-M. The representation obtained by dereferencing M is m. A timemap, $\mathbb{T}$, for an R is a set of M, which can be empty. A composite memento, $\mathbb{C}$, is a root URI-M, M_0, and zero or more embedded URI-Ms, $M_{i,j}$, one for each R_i (images, css, etc.).

The recomposition process is represented by algorithm (1). Three parameters are required: R is a URI-R, t is a target datetime, and $\mathcal{H}$ is a memento selection heuristic. The algorithm states that a com-

Table 2: Functions

Function	Definition
$\mathcal{E}(m)$	Set of Rs embedded in m.
$\mathcal{H}(R,t)$	Heuristic providing the best M for R at t.
$\mathcal{G}(M)$	Representation of M; generally retrieved via HTTP GET.
$\mathcal{C}(R,t,\mathcal{H})$	Composite memento for R.

†See RFC 2616 [6].

posite memento is the set of recursively selected URI-Ms, one per URI-R, starting with URI-R, R, for target-datetime t, using heuristic $\mathcal{H}$. Heuristic $\mathcal{H}$ is a function that returns the best[1] M for R at target datetime t. Thus, $\mathcal{G}(\mathcal{H}(R,t))$ is the best memento for R at t under $\mathcal{H}$. The rest of the algorithm, $\{\mathcal{C}(R_i,t,\mathcal{H}) \mid R_i \in \mathcal{E}(\mathcal{G}(\mathcal{H}(R,t)))\}$, is the set of mementos for the URI-Rs embedded in R, applied recursively.

$$\mathcal{C}(R,t,\mathcal{H}) = \mathcal{G}(\mathcal{H}(R,t)) \cup \{\mathcal{C}(R_i,t,\mathcal{H}) \mid R_i \in \mathcal{E}(\mathcal{G}(\mathcal{H}(R,t)))\} \quad (1)$$

Given algorithm (1), the composite memento, $\mathbb{C}$, for root URI-R, R_0, for target datetime t, under heuristic $\mathcal{H}$ is shown in (2).

$$\mathbb{C} = \mathcal{C}(R_0,t,\mathcal{H}) \quad (2)$$

2.3 Temporal Spread

Temporal spread is the difference between the earliest and latest Memento-Datetimes in a composite memento. Consider again the December 9, 2004 composite memento for the wunderground.org page, which is shown again in figure 3. It comprises a root memento[2] and 128 embedded resources, a sample of which are shown in table 3. The earliest capture occurred 1.8 months before the root was captured. The last capture occurred *8.1 years* after the root was captured. The temporal spread is 8.3 years, the mean delta is 1.8 years and the standard deviation is 2.9 years. Thirty four of the embedded URI-Rs are not archived (e.g., the banner ad labeled missing in figure 3) and one is archived (a timemap exists) but the desired embedded memento is not availble.

2.4 Temporal Coherence

We define an embedded memento to be temporally coherent with respect to a root memento when it can be shown that the embedded memento's representation existed at the time the root memento was captured. So, which of the embedded mementos listed in table 3 are temporally coherent? The answer to this question requires evaluating the relationship between the root memento and the embedded memento. This relationship can be one of many patterns, described in section 3.

3. TEMPORAL COHERENCE PATTERNS

The temporal coherence patterns defined below occurs are based on the relationship between Memento-Datetime, Last-Modified datetime, and entity body content of a root and one or two embedded mementos. Symbols for these attributes are defined in table 4. The relationships are illustrated in charts like the one in figure 4. Each diamond represents a memento. The red hollow diamond is the best

[1]The definition of *best* depends on the heuristic. A common definition is least difference between capture datetime and target datetime. RFC 7089 [13] does not specify a memento heuristic.

[2]`http://web.archive.org/web/20041209190926/http://www.wunderground.org/cgi-bin/findWeather/getForecast?query=50593`

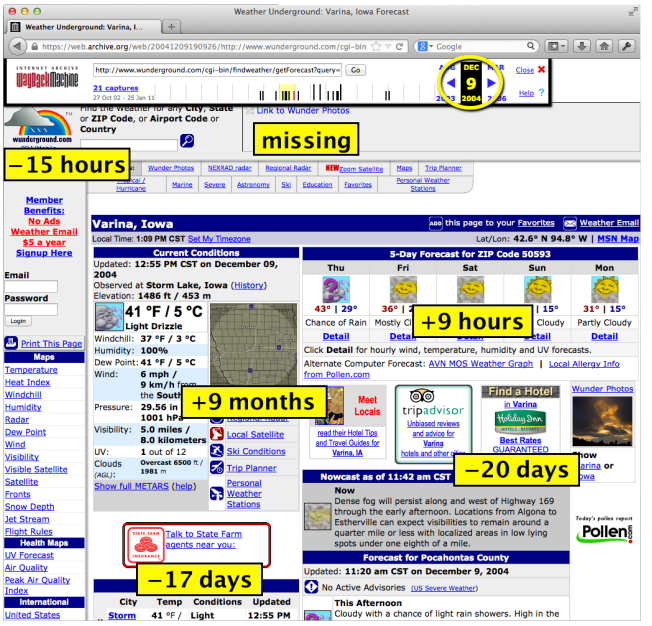

Figure 3: wunderground.org forecast page, Dec. 9, 2004

Table 4: Memento Attributes

Term	Definition
$\mathcal{T}_0$	*Memento-Datetime* of root memento m_0.
$\mathcal{L}_0$	*Last-Modified*† datetime of root memento m_0.
$\mathcal{T}_{i,j}$	*Memento-Datetime* of embedded memento $m_{i,j}$.
$\mathcal{L}_{i,j}$	*Last-Modified*† datetime of memento embedded $m_{i,j}$
$\mathcal{B}_{i,j}$	Entity body† of memento embedded $m_{i,j}$.

†See RFC 2616 [6].

memento for the root; solid diamonds are mementos for an embedded resource. Blue objects affect the coherence state; gray objects do not. The long horizontal black line is a time line, with earlier datetimes to the left as usual. Mementos and Last-Modified datetimes are plotted relative to each other. The time line is not scaled; the distance between datetimes is not meaningful. Last-Modified datetimes are associated with mementos by a line pointing from the memento diamond to $\mathcal{L}$. Undefined data is denoted with an up arrow (↑). Likewise, a down arrow (↓) asserts that an attribute is or must be defined. (Arrow notation adopted from Kieffer et al. [7].)

3.1 Temporal Coherence States

Each pattern is categorized into one of five temporal coherence states:

Prima Facie Coherent (C) The embedded memento existed in its archived state at the time the root memento was captured.

Prima Facie Violative (V) The embedded memento did not exist

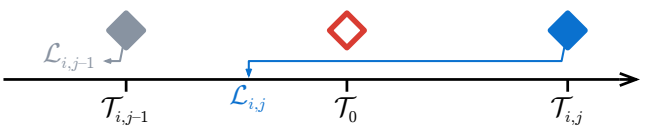

Figure 4: Sample Pattern Chart

Table 3: Embedded Memento Capture Datetimes (15 of 128)

URI	Memento-Datetime	Delta
http://ads.wunderground.com/ads/images/wu-9.jpg	2004-10-16 03:40:53Z	-1.8 months
http://icons.wunderground.com/graphics/smalllogo2.gif	2004-11-22 05:46:03Z	-18 days
http://ads.wunderground.com/ads/images/Statefarm-sfcom001_120x30.gif	2004-11-22 05:46:32Z	-18 days
http://icons.wunderground.com/ads/images/Davi-00009-vp2_125x125.gif	2004-12-08 07:01:08Z	-2 days
http://icons.wunderground.com/graphics/smash/wunderTransparent.gif	2004-12-09 04:36:14Z	-15 hours
http://www.wunderground.com/cgi-bin/findweather/getForecast?query=50593	2004-12-09 19:29:26Z	– *root*
http://icons.wunderground.com/graphics/conds/cloudy.GIF	2004-12-10 04:48:55Z	+21 minutes
http://icons.wunderground.com/graphics/conds/mostlycloudy.GIF	2004-12-10 04:48:55Z	+10 hours
http://icons.wunderground.com/graphics/conds/rait.GIF	2004-12-12 14:54:01Z	+3 days
http://icons.wunderground.com/ads/images/TripAdvisor-Blinky.gif	2005-01-27 02:58:30Z	+1.6 months
http://banners.wunderground.com/cgi-bin/statefarmbanner?zip=50593&width=150	2006-03-26 03:29:00Z	+1.3 years
http://www.valueclick.com/system/files/coupon-mountain-slider.png?1310511429	2011-07-13 09:08:00Z	+6.6 years
http://z1.adserver.com/system/files/logo.gif	2013-01-13 06:09:14Z	+8.1 years
http://icons.wunderground.com/data/wximagenew/d/d70dave/0-thumb.jpg	Not Archived	
http://pagead2.googlesyndication.com/pagead/show_ads.js	Missing Memento	

in its archived state at the time the root memento was captured.

Possibly Coherent (PC) The embedded memento could have existed in its archive state at the time the root memento was captured.

Probably Violative (PV) The embedded memento probably did not exist in its archived state at the time the root memento was captured.

Coherence Undefined (CU) There is not enough information to determine coherence state.

3.2 Pattern Groups

The patterns are presented in four groups. One-memento patterns consider a single memento for an embedded resource. Two-memento patterns consider two mementos, one captured on or before the root and one captured after. Content patterns are two-memento patterns that also consider memento content, not just capture and modification datetimes. The other patterns group includes special cases such as a root memento that has no embedded resources.

3.2.1 One-Memento Patterns

One-memento patterns consider one root URI-R, R_0, and one embedded URI-R, R_i, and their corresponding mementos, m_0 and $m_{i,j}$ respectively. If $\mathcal{T}_{i,j} < \mathcal{T}_0$ for all $m_{i,j}$, then a *left*[3] pattern exists and only the newest memento, $m_{i,n}$, must be considered. If $\mathcal{T}_{i,j} \geq \mathcal{T}_0$ for all $m_{i,j}$, then a *right* pattern exists and only the oldest memento, $m_{i,1}$, must be considered.

1RB: Right Bracket

The right pattern depicted in figure 5(a), and specified by predicate (3), represents an embedded memento, $m_{i,1}$, which was captured after the root memento, m_0, was captured, but was modified before m_0 was captured The embedded memento's Memento-Datetime is after the root's and its Last-Modified datetime is on or before the root's Memento-Datetime. The embedded memento's Last-Modified datetime and its Memento-Datetime bracket the root's Memento-Datetime. Therefore, the embedded memento existed in its archived state at the time the root memento was captured. Pattern 1RB is prima facie coherent.

$$\mathcal{L}_{i,1}\downarrow \wedge (\mathcal{L}_{i,1} \leq \mathcal{T}_0 < \mathcal{T}_{i,1}) \implies \text{C} \tag{3}$$

[3] *Left* and *right* come from the appearance of the charts.

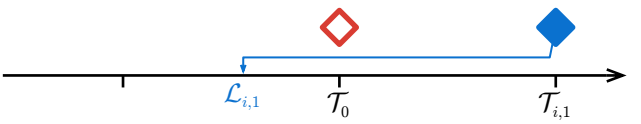

(a) Right Bracket (1RB) ⇒ C

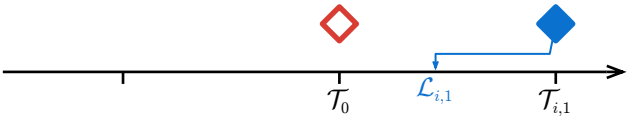

(b) Right Newer Last-Modified (1RN) ⇒ V

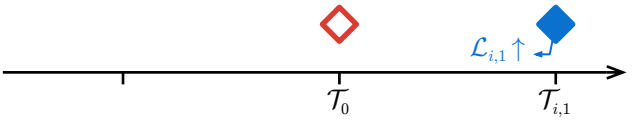

(c) Right Undefined Last-Modified (1RU) ⇒ PV

Figure 5: Right-Sided Patterns

1RN: Right Newer Last-Modified

The right pattern depicted in figure 5(b), and specified by predicate (4), represents an embedded memento, $m_{i,1}$, which was both modified and captured after the root memento, m_0, was captured. The embedded memento's Memento-Datetime and Last-Modified datetime are both later than the root's Memento-Datetime. This evidence indicates that the embedded memento was modified after the root memento; therefore, the embedded memento did not exist in its archived state at the time the root memento was captured. Pattern 1RN is prima facie violative.

$$\mathcal{L}_{i,1}\downarrow \wedge (\mathcal{T}_0 < \mathcal{L}_{i,1} \leq \mathcal{T}_{i,1}) \implies \text{V} \tag{4}$$

1RU: Right Undefined Last-Modified

The right pattern depicted in figure 5(c), and specified by predicate (5), represents an embedded memento, $m_{i,1}$, which was captured after the root memento, m_0, and does not have a Last-Modified datetime. The embedded memento's Memento-Datetime is after the root's, like the 1RN pattern, but in this case the Last-Modified datetime is undefined. This evidence alone does not allow determination of the embedded memento's state at the time the root me-

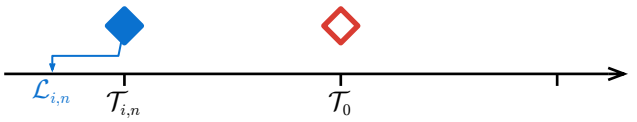

(a) Left Last-Modified (1LL) $\Rightarrow$ PC

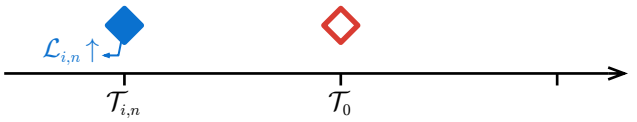

(b) Left Undefined Last-Modified (1LU) $\Rightarrow$ PV

Figure 6: Left-Sided Patterns

mento was captured. However, since the embedded URI-R is referenced by the root memento, it is likely that a representation existed at the time the root was captured. Our experience is that a missing Last-Modified datetime header normally indicates a dynamically-generated resource generated on demand (explanation in Appendix B). Thus, pattern 1RU is probably violative.

$$\mathcal{L}_{i,1}\uparrow \wedge (\mathcal{T}_0 < \mathcal{T}_{i,1}) \implies \text{PV} \tag{5}$$

1LL: Left Last-Modified

The left pattern depicted in figure 6(a), and specified by predicate (6), represents an embedded memento, $m_{i,n}$, which was captured before the root memento, m_0, and has a Last-Modified datetime. Because the embedded memento's Memento-Datetime is before the root's, the Last-Modified datetime does not directly affect coherence state. However, the existence of Last-Modified indicates that the embedded memento was probably not dynamically generated. Therefore, pattern 1LL is probably coherent.

$$\mathcal{L}_{i,n}\downarrow \wedge (\mathcal{T}_{i,n} < \mathcal{T}_0) \implies \text{PC} \tag{6}$$

1LU: Left Undefined Last-Modified

The left pattern depicted in figure 6(b), and specified by predicate (7), represents an embedded memento, $m_{i,n}$, which was captured before the root memento, m_0, and does not have a Last-Modified datetime. The embedded memento's Memento-Datetime is before the root's, which means the lack of a Last-Modified datetime does not directly affect coherence state. However, the lack of Last-Modified implies that the embedded resource was probably dynamically generated. Therefore, pattern 1LU is probably violative.

$$\mathcal{L}_{i,n}\uparrow \wedge (\mathcal{T}_{i,n} < \mathcal{T}_0) \implies \text{PV} \tag{7}$$

1EQ: Simultaneous Capture

The pattern depicted in figure 7, and specified by predicate (8), represents an embedded memento captured simultaneously with the root[4]. Because the embedded memento's Memento-Datetime equals the root's, pattern 1EQ is prima facie coherent.

$$\mathcal{T}_0 = \mathcal{T}_{i,j} \implies \text{C} \tag{8}$$

[4]Exact simultaneity is improbable. However, Memento-Datetime precision is one second and a web browser is unlikely to download multiple copies within one second. Therefore, the captures are effectively simultaneous.

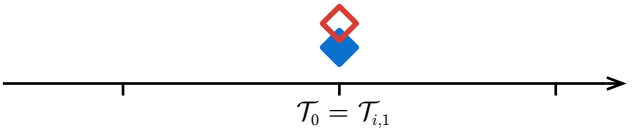

Figure 7: Simultaneous Capture $\Rightarrow$ C

3.2.2 *Two-Memento Patterns*

Two-memento patterns consider one root URI-R, R_0, with a single memento, m_0, and one embedded URI-R, R_i, with two consecutive mementos, $m_{i,j-1}$ and $m_{i,j}$. The consecutive mementos are selected such that $\mathcal{T}_{i,j-1} < \mathcal{T}_0 < \mathcal{T}_{i,j}$. All three two-memento patterns share this characteristic.

The two-memento patterns are closely related to similar right (1RB, 1RN, and 1RU) patterns. The primary difference is the addition of a left memento to the two-memento patterns. The two-memento patterns are named 2B, 2N, and 2U after their one-memento counterparts and are depicted in figure 8. Note that all the left mementos are shown in gray, which denotes that they do not affect coherence state.

2B: Two-Memento Bracket

The two-memento pattern depicted in figure 8(a), and specified by predicate (9), represents a pair of mementos for an embedded resource, $m_{i,j-1}$ and $m_{i,j}$, with $m_{i,j-1}$ captured before the root memento and $m_{i,j}$ captured after the root memento. $m_{i,j}$ has a Last-Modified datetime that is on or before the root's capture time. Thus, $m_{i,j}$ is under the same conditions in this pattern as it is in 1RB. This can be seen by comparing (9) with (3). Therefore, the embedded memento existed in its archived state at the time the root memento was captured and Pattern 2B is prima facie coherent. (Note: Because $m_{i,j}$ brackets m_0, $m_{i,j-1}$ does not affect temporal coherence state.)

$$\mathcal{L}_{i,j}\downarrow \wedge (\mathcal{T}_{i,j-1} < \mathcal{L}_{i,j} < \mathcal{T}_0 < \mathcal{T}_{i,1}) \implies \text{C} \tag{9}$$

2N: Two-Memento Newer Last-Modified

The two-memento pattern depicted in figure 8(b), and specified by predicate (10), represents a pair of mementos for an embedded resource, $m_{i,j-1}$ and $m_{i,j}$, the same capture timing as in pattern 2B. This pattern differs from 2B in that $m_{i,j}$ has a Last-Modified datetime that after root's capture time. Thus, $m_{i,j}$ is under the same conditions in this pattern as it is in 1RN. This can be seen by comparing (10) with (4). Therefore, like pattern 2RN, pattern 2N is prima facie violative. (Note: Because the state of $m_{i,j-1}$ is unknown from $\mathcal{T}_{i,j-1}$ to $\mathcal{L}_{i,j}$, it does not affect temporal coherence state.)

$$\mathcal{L}_{i,1}\downarrow \wedge (\mathcal{T}_{i,j-1} < \mathcal{T}_0 < \mathcal{L}_{i,1} < \mathcal{T}_{i,1}) \implies \text{V} \tag{10}$$

2U: Two-Memento Undefined Last-Modified

The two-memento pattern depicted in figure 8(c), and specified by predicate (11), represents a pair of mementos for an embedded resource, $m_{i,j-1}$ and $m_{i,j}$, the same capture timing as in patterns 2B and 2N. This pattern differs from 2B and 2N in that $m_{i,j}$ has an undefined Last-Modified datetime. Thus, $m_{i,j}$ is under the same conditions in this pattern as it is in 1RU. This can be seen by comparing (11) with (5). Therefore, like pattern 1RU, pattern 2U is probably violative.

$$\mathcal{L}_{i,1}\uparrow \wedge (\mathcal{T}_{i,j-1} < \mathcal{T}_0 < \mathcal{T}_{i,1}) \implies \text{PV} \tag{11}$$

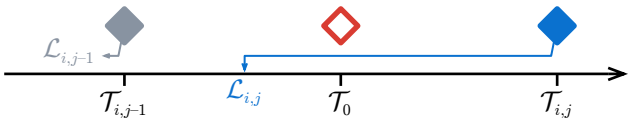

(a) Bracket (2B) ⇒ C

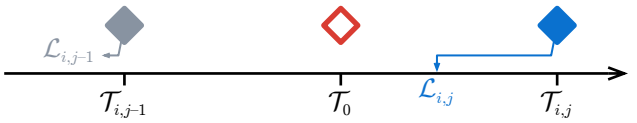

(b) Newer Last-Modified (2N) ⇒ V

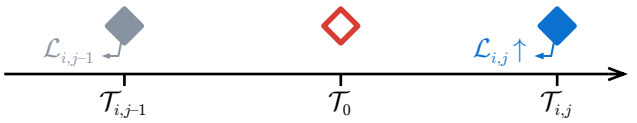

(c) Undefined Last-Modified (2U) ⇒ PV

Figure 8: Two-Memento Patterns

3.2.3 *Content Patterns*

Even when Last-Modified datetime is invalid or unavailable, when two mementos are available, additional evidence is available: the mementos' content, and more importantly, their similarity or lack thereof.

Content patterns are extensions to the two-memento patterns that add memento content to the determination of coherence state. Content patterns require more computing resources and time than the one- and two-memento patterns because two memento entity bodies must be retrieved for each embedded URI-R. This additional cost may render content patterns unsuitable for casual archive use or in restricted bandwidth conditions.

The Content patterns are depicted in figure 9. Each of the three charts represents a class of patterns. These pattern classes are extensions of the two-memento patterns defined in 3.2.2, differing in how memento content is evaluated and how it affects coherence state. These differences are represented by the asterisk operator ($*$), which in turn represents an evaluation function. Evaluation functions provide evidence about the archival state of the embedded memento, $m_{i,j-1}$, at the time the root memento, m_0, was captured. There are three evaluation functions:

Equals ($=$) $\mathcal{B}_{m_{i,j-1}} = \mathcal{B}_{m_{i,j}}$. This function returns true if the bodies of $m_{i,j-1}$ and $m_{i,j}$ are bit-for-bit equal. A true value indicates that $m_{i,j-1}$ probably existed in its archived state at the time m_0 was captured.

Similar ($\sim$) $\mathcal{B}_{m_{i,j-1}} \sim \mathcal{B}_{m_{i,j}}$. This function returns true if the bodies of $m_{i,j-1}$ and $m_{i,j}$ are substantially similar. A true value asserts with high confidence that $m_{i,j-1}$ existed in its archived state at the time the m_0 was captured.

Not Similar ($\nsim$) $\mathcal{B}_{m_{i,j-1}} \nsim \mathcal{B}_{m_{i,j}}$. This function returns true if the bodies of $m_{i,j-1}$ are not substantially similar. Note, for mementos, $m_i \nsim m_j \implies m_i \neq m_j$.

The reason for both equality and similarity is that many web archives treat text resources (e.g., HTML) and binary resources (e.g., images) differently. Archive metadata is frequently added to text resources, while binary resources are not changed. It makes sense to first check equality, and fall back to similarity if the equality check fails. It should be noted that while the definition equality is universal, the definition of similar will vary by application and user need.

All of the equality and similarity patterns result in the Prima Facie Coherent temporal coherence state. This raises the question of whether or not all six patterns are required. Although a single equality pattern and single similarity pattern are sufficient, we have chosen to retain all six patterns for the present.

2EB: Content Equal Bracket

The content pattern depicted in figure 9(a), and specified by predicate (12), is the same as pattern 2B plus the determination that the two embedded mementos have bit-for-bit equality. Bit-for-bit equality is an overarching condition, which when combined with the Last-Modified datetime evidence described in patterns 1RB and 2B, provides strong evidence the embedded memento existed in its archived state at the time the root memento was captured. Thus, pattern 2EB is prima facie coherent.

$$\mathcal{L}_{i,j}\downarrow \wedge (\mathcal{T}_{i,j-1} < \mathcal{L}_{i,j} < \mathcal{T}_0 < \mathcal{T}_{i,1}) \wedge (m_{i,j-1} = m_{i,j}) \implies \text{C} \quad (12)$$

2EN: Content Equal Newer Last-Modified

The content pattern depicted in figure 9(b), and specified by predicate (13), is the same as pattern 2N plus the determination that the two embedded mementos have bit-for-bit equality. Bit-for-bit equality is an overarching condition which provides evidence that the embedded memento existed in its archived state at the time the root memento was captured; it overrides the Last-Modified datetime evidence described in related patterns 1RN and 2N. Thus, pattern 2EN is prima facie coherent.

$$\mathcal{L}_{i,1}\downarrow \wedge (\mathcal{T}_{i,j-1} < \mathcal{T}_0 < \mathcal{L}_{i,1} < \mathcal{T}_{i,1}) \wedge (m_{i,j-1} = m_{i,j}) \implies \text{C} \quad (13)$$

2EU: Content Equal Undefined Last-Modified

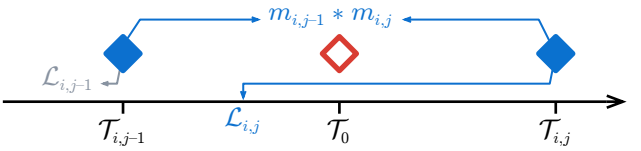

(a) Bracket w/Evaluation (2EB, 2SB, 2NB) ⇒ (C,C,C)

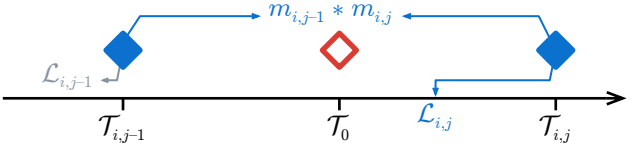

(b) Newer Last-Modified w/Content (2EN, 2SN, 2NN) ⇒ (C,C,C)

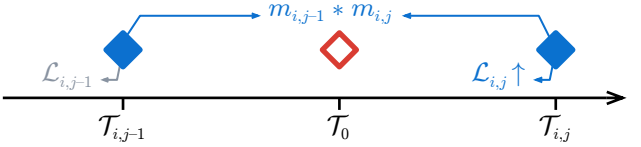

(c) Undef. Last-Modified w/Content (2EU, 2SU, 2NU) ⇒ (C,C,V)

Figure 9: Content Patterns

The content pattern depicted in figure 9(c), and specified by predicate (14), is the same as pattern 2U plus the determination that the two embedded mementos have bit-for-bit equality. Bit-for-bit equality is an overarching condition which provides evidence that the embedded memento existed in its archived state at the time the root memento was captured; it overrides the undefined Last-Modified datetime evidence described in related patterns 1RU and 2U. Thus, pattern 2EU is prima facie coherent.

$$\mathcal{L}_{i,1}\uparrow \wedge (\mathcal{T}_{i,j-1} < \mathcal{T}_0 < \mathcal{T}_{i,1}) \wedge (m_{i,j-1} = m_{i,j}) \implies \text{C} \qquad (14)$$

2SB: Content Similar Bracket
The content pattern depicted in figure 9(a), and specified by predicate (15), is a weaker form of pattern 2EB. In place of bit-for-bit equality, a similarity measure it used. Like equality, sufficient similarity is an overarching condition and provides evidence that the embedded memento existed in its archived state at the time the root memento was captured; it overrides the Last-Modified datetime evidence described in related patterns 1RB and 2B. Thus, pattern 2SB is prima facie coherent.

$$\mathcal{L}_{i,j}\downarrow \wedge (\mathcal{T}_{i,j-1} < \mathcal{L}_{i,j} < \mathcal{T}_0 < \mathcal{T}_{i,1}) \wedge (m_{i,j-1} \sim m_{i,j}) \implies \text{C} \quad (15)$$

2SN: Content Similar Newer Last-Modified
The content pattern depicted in figure 9(b), and specified by predicate (16), a weaker form of pattern 2EN. Like pattern 2SB, a similarity measure is used in place of bit-for-bit equality. Like equality, sufficient similarity is an overarching condition and provides evidence that the embedded memento existed in its archived state at the time the root memento was captured; it overrides the Last-Modified datetime evidence described in related patterns 1RN and 2N. Thus, pattern 2EN is prima facie coherent.

$$\mathcal{L}_{i,1}\downarrow \wedge (\mathcal{T}_{i,j-1} < \mathcal{T}_0 < \mathcal{L}_{i,1} < \mathcal{T}_{i,1}) \wedge (m_{i,j-1} \sim m_{i,j}) \implies \text{C} \quad (16)$$

2SU: Content Similar Undefined Last-Modified
The content pattern depicted in figure 9(c), and specified by predicate (17), a weaker form of pattern 2EU. Like patterns 2SB and 2SN, a similarity measure is used in place of bit-for-bit equality. Like equality, sufficient similarity is an overarching condition and provides evidence that the embedded memento existed in its archived state at the time the root memento was captured. This similarity overrides the Last-Modified datetime evidence described in related patterns 1RU and 2U.

$$\mathcal{L}_{i,1}\uparrow \wedge (\mathcal{T}_{i,j-1} < \mathcal{T}_0 < \mathcal{T}_{i,1}) \wedge (m_{i,j-1} \sim m_{i,j}) \implies \text{C} \qquad (17)$$

2NB: Content Not Similar Bracket
The content pattern depicted in figure 9(a), and specified by predicate (18), is the same as pattern 2B plus the determination that the two embedded mementos have different content. Unlike the substantial similarity of pattern 2EB, the lack of similarity is not an overarching condition and does not add significant evidence that the embedded memento existed in its archived state at the time the root memento was captured. Pattern 2NB is therefore equivalent to pattern 2B, which is prima facie coherent.

$$\mathcal{L}_{i,j}\downarrow \wedge (\mathcal{T}_{i,j-1} < \mathcal{L}_{i,j} < \mathcal{T}_0 < \mathcal{T}_{i,1}) \wedge (m_{i,j-1} \nsim m_{i,j}) \implies \text{C} \quad (18)$$

2NN: Content Not Similar Newer Last-Modified
The content pattern depicted in figure 9(b), and specified by predicate (19), is the same as pattern 2N plus the determination that the two embedded mementos have different content. Unlike the substantial similarity of pattern 2EN, the lack of similarity is not an overarching condition and does not add significant evidence that the embedded memento existed in its archived state at the time the root memento was captured. Pattern 2NN is therefore equivalent to pattern 2N, which is prima facie violative.

$$\mathcal{L}_{i,1}\downarrow \wedge (\mathcal{T}_{i,j-1} < \mathcal{T}_0 < \mathcal{L}_{i,1} < \mathcal{T}_{i,1}) \wedge (m_{i,j-1} \nsim m_{i,j}) \implies \text{V} \quad (19)$$

2NU: Content Not Similar Undefined Last-Modified
The content pattern depicted in figure 9(c), and specified by predicate (20), is the same as pattern 2U plus the determination that the two embedded mementos have different content. Unlike the substantial similarity of pattern 2EU, the lack of similarity is not an overarching condition and does not add significant evidence that the embedded memento existed in its archived state at the time the root memento was captured. Pattern 2NU is therefore equivalent to pattern 2U, which is probably violative.

$$\mathcal{L}_{i,1}\uparrow \wedge (\mathcal{T}_{i,j-1} < \mathcal{T}_0 < \mathcal{T}_{i,1}) \wedge (m_{i,j-1} \nsim m_{i,j}) \implies \text{PV} \qquad (20)$$

3.2.4 Other Patterns

There are several patterns that consider one root URI-R, R_0, as do the one- and two-memento patterns, but do not consider embedded mementos. The lack of embedded mementos means these patterns are either prima facie coherent or have undefined coherence state.

0NE: No Embedded URIs
The pattern depicted in figure 10(a) represents a root memento, m_0, which has no embedded URIs (e.g., images, plain text). These root mementos have inherent temporal coherence.

0NA: Not Archived
The pattern depicted in figure 10(b) represents a root memento, m_0, and an embedded URI-R, R_i, which is not archived (or the mementos are not available). Temporal coherence state for this pattern is undefined.

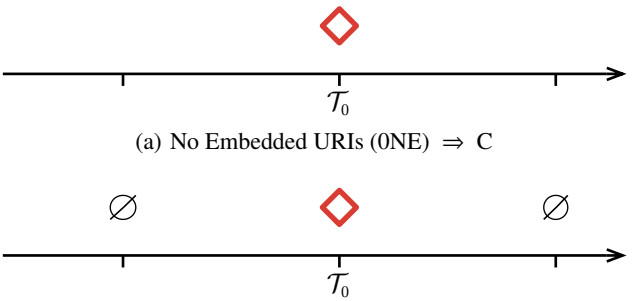

(a) No Embedded URIs (0NE) $\Rightarrow$ C

(b) Not Archived (0NA) $\Rightarrow$ UC

Figure 10: Other Patterns

4. WEB DATA IS IMPERFECT

The pattern descriptions in subsections 3.2.1–3.2.3 do not address the vast amounts of invalid and badly structured data present on the Web; and, consequently part of Web archive holdings. This section describes a few of the imperfections we routinely encounter.

4.1 Invalid Datetimes

Datetimes in particular are critical to the identification of patterns and proper classification of coherence state. In our work to date,

Table 5: Datetime Corrections

Symptom	Description and Corrective Action
2-Digit year	This is symptomatic of correctly functioning pre-year 2000 software. Corrected by adding 1900 to the year.
1-Digit year	This Y2K bug produced the correct 2-digit year prior to 2000, but fails from 2000 on. We have seen this error on datetimes as late as 2003. Corrected by adding 2000 to the year.
3-Digit year	Another Y2K bug caused by subtracting 1900 to produce a 2-digit year. We have seen this error on datetimes as late as 2002. Corrected by adding 1900 to the year.
Non-GMT	Local time zones are common in mementos captured before 2000. Corrected by converting to GMT.
Non-English	Most datetime parsers only handle English month and weekday names, yet other languages crop up in datetime strings. The most common is French. Corrected using a lookup table to translate to English.
Missing zeros	Missing leading zeros are common. Most common is day of month. Time fields are also affected.
Extra zeros	Although not as common as missing leading zeros, extra leading zeros are found. Corrected by ignoring extra leading zeros.

every effort is made to *decode* invalid datetimes. Table 5 lists the major corrective actions currently used. If a datetime cannot be parsed, it is treated as undefined.

4.2 Last-Modified Datetimes

Clausen [4] demonstrated that Last-Modified datetimes are a reliable indicator of entity body change. Therefore, we treat most Last-Modified datetimes as valid. Clearly incorrect Last-Modified datetimes are treated as undefined (for example, when the original resource Last-Modified datetime is greater than the Memento-Datetime). The two basic cases are (1) when the original resource server sets Last-Modified to the current datetime on every response, and (2) when the endity body is modified but original resource Last-Modified is not. Both render attempts to determine coherence ineffective; fortunately, both are rare. Much more common is the original resources server omitting Last-Modified altogether (patterns 1RU, 1LU, and 2U). The content patterns defined in 3.2.3 (2EU, 2SU, and 2NU in particular) address the ommission using entity body equality and similarity.

It must also be noted that there are pathological cases caused by misconfigured servers and buggy software, where changes to the memento content and changes to Last-Modified datetime are unrelated.

4.3 Missing Mementos

Many of the patterns rely on retrieving memento HTTP headers and content. However, this is not always possible. The causes are not important, but the consequences must be considered. For example, two consecutive mementos, $m_{i,j-1}$ and $m_{i,j}$, would normally result in a two-memento pattern match. But what if $m_{i,j-1}$ cannot be retrieved? Should an $m_{i,j-2}$ be used as substitute? Should a one-memento pattern be used instead? These are open questions, the answers to which may depend on circumstances and the user priorities.

4.4 Memento-Datetime Collisions

Occasionally, an archive will have multiple mementos with the same Memento-Datetime. There are many causes; three are briefly described here. First, several archives can capture the same original resource at the same time. If a client later uses multiple archives simultaneously, multiple mementos with the same Memento-Datetime will be found. Second, a single capture is available from multiple archives because the archives share holdings (e.g., Archive-It! and the Internet Archive). In both of these cases, the content is likely to be identical; however, metadata differences may exist (e.g., one archive captures Last-Modifed datetime and the other does not). It is therefore possible for more than one pattern to be matched. In this case, we suggest using the least favorable coherence state. However, like the missing mementos issue, the best response to this condition may depend on user priorities.

A third scenario that leads to Memento-Datetime collisions is typified by Wikipedia edit wars [1]. In this case, changes occur less than one second apart. Because Memento-Datetime inherits one-second resolution from HTTP, the resulting mementos can have the same Memento-Datetime.

4.5 Redirection to the Live Web

The Internet Archive Wayback Machine makes every effort to fulfill requests for embedded mementos. In some cases this means redirecting to a Live Web resource. In our research, we detect this redirection and treat the URI-R as if it has no mementos.

5. CONCLUSION

We have introduced temporal coherence patterns of embedded mementos and four coherence states: prima facie coherent, possibly coherent, probably violative, and prima facie violative. The patterns and resulting states, are summarized in table 6, Together, the patterns and states provide a framework in which to examine the temporal coherence of embedded mementos.

6. REFERENCES

[1] Wikipedia:edit warring. `http://en.wikipedia.org/wiki/Wikipedia:Edit_war`.

[2] Myriam Ben Saad and Stéphane Gançarski. Archiving the Web using page changes patterns: a case study. In Proceedings of the 11th annual international ACM/IEEE joint conference on Digital libraries, JCDL '11, pages 113–122, 2011.

[3] Myriam Ben Saad, Zeynep Pehlivan, and Stéphane Gançarski. Coherence-oriented crawling and navigation using patterns for web archives. In Proceedings of the 15th international conference on Theory and practice of digital libraries: research and advanced technology for digital libraries, TPDL'11, pages 421–433, 2011.

[4] Lars Clausen. Concerning ETags and datestamps. In 4th International Web Archiving Workshop (IWAW'04), 2004.

[5] Dimitar Denev, Arturas Mazeika, Marc Spaniol, and Gerhard Weikum. SHARC: Framework for quality-conscious web archiving. In Proceedings of the VLDB Endowment, volume 2, pages 586–597, August 2009.

[6] R. T. Fielding, J. Gettys, J. C. Mogul, H. F. Nielsen, L. Masinter, P. Leach, and T. Berners-Lee. Hypertext transfer protocol—HTTP/1.1 (IETF RFC 2616), June 1999. `http://www.ietf.org/rfc/rfc2616.txt`.

Table 6: Coherence Patterns Summary

Abbr.	Description	Condition	Coherence State
One-Memento Patterns			
1RB	Right Bracket	$\mathcal{L}_{i,1}\downarrow \wedge (\mathcal{L}_{i,1} \leq \mathcal{T}_0 < \mathcal{T}_{i,1})$	Prima Facie Coherent
1RN	Right Newer Last-Modified	$\mathcal{L}_{i,1}\downarrow \wedge (\mathcal{T}_0 < \mathcal{L}_{i,1} \leq \mathcal{T}_{i,1})$	Prima Facie Violative
1RU	Right Undefined Last-Modified	$\mathcal{L}_{i,1}\uparrow \wedge (\mathcal{T}_0 < \mathcal{T}_{i,1})$	Probably Violative
1LL	Left Last-Modified	$\mathcal{L}_{i,n}\downarrow \wedge (\mathcal{T}_{i,n} < \mathcal{T}_0)$	Possibly Coherent
1LU	Left Undefined Last-Modified	$\mathcal{L}_{i,n}\uparrow \wedge (\mathcal{T}_{i,n} < \mathcal{T}_0)$	Probably Violative
1EQ	Simultaneous Capture	$\mathcal{T}_0 = \mathcal{T}_{i,j}$	Prima Facie Coherent
Two-Memento Patterns			
2B	Bracket	$\mathcal{L}_{i,j}\downarrow \wedge (\mathcal{T}_{i,j-1} < \mathcal{L}_{i,j} < \mathcal{T}_0 < \mathcal{T}_{i,1})$	Prima Facie Coherent
2N	Newer Last-Modified	$\mathcal{L}_{i,1}\downarrow \wedge (\mathcal{T}_{i,j-1} < \mathcal{T}_0 < \mathcal{L}_{i,1} < \mathcal{T}_{i,1})$	Prima Facie Violative
2U	Undefined Last-Modified	$\mathcal{L}_{i,1}\uparrow \wedge (\mathcal{T}_{i,j-1} < \mathcal{T}_0 < \mathcal{T}_{i,1})$	Probably Violative
Content Patterns			
2EB	Equal Bracket	$\mathcal{L}_{i,j}\downarrow \wedge (\mathcal{T}_{i,j-1} < \mathcal{L}_{i,j} < \mathcal{T}_0 < \mathcal{T}_{i,1}) \wedge (m_{i,j-1} = m_{i,j})$	Prima Facie Coherent
2EN	Equal Newer Last-Modified	$\mathcal{L}_{i,1}\downarrow \wedge (\mathcal{T}_{i,j-1} < \mathcal{T}_0 < \mathcal{L}_{i,1} < \mathcal{T}_{i,1}) \wedge (m_{i,j-1} = m_{i,j})$	Prima Facie Coherent
2EU	Equal Undefined Last-Modified	$\mathcal{L}_{i,1}\uparrow \wedge (\mathcal{T}_{i,j-1} < \mathcal{T}_0 < \mathcal{T}_{i,1}) \wedge (m_{i,j-1} = m_{i,j})$	Prima Facie Coherent
2SB	Similar Bracket	$\mathcal{L}_{i,j}\downarrow \wedge (\mathcal{T}_{i,j-1} < \mathcal{L}_{i,j} < \mathcal{T}_0 < \mathcal{T}_{i,1}) \wedge (m_{i,j-1} \sim m_{i,j})$	Prima Facie Coherent
2SN	Similar Newer Last-Modified	$\mathcal{L}_{i,1}\downarrow \wedge (\mathcal{T}_{i,j-1} < \mathcal{T}_0 < \mathcal{L}_{i,1} < \mathcal{T}_{i,1}) \wedge (m_{i,j-1} \sim m_{i,j})$	Prima Facie Coherent
2SU	Similar Undefined Last-Modified	$\mathcal{L}_{i,1}\uparrow \wedge (\mathcal{T}_{i,j-1} < \mathcal{T}_0 < \mathcal{T}_{i,1}) \wedge (m_{i,j-1} \sim m_{i,j})$	Prima Facie Coherent
2NB	Not Similar Bracket	$\mathcal{L}_{i,j}\downarrow \wedge (\mathcal{T}_{i,j-1} < \mathcal{L}_{i,j} < \mathcal{T}_0 < \mathcal{T}_{i,1}) \wedge (m_{i,j-1} \nsim m_{i,j})$	Prima Facie Coherent
2NN	Not Similar Newer Last-Modified	$\mathcal{L}_{i,1}\downarrow \wedge (\mathcal{T}_{i,j-1} < \mathcal{T}_0 < \mathcal{L}_{i,1} < \mathcal{T}_{i,1}) \wedge (m_{i,j-1} \nsim m_{i,j})$	Prima Facie Violative
2NU	Not Similar Undefined Last-Modified	$\mathcal{L}_{i,1}\uparrow \wedge (\mathcal{T}_{i,j-1} < \mathcal{T}_0 < \mathcal{T}_{i,1}) \wedge (m_{i,j-1} \nsim m_{i,j})$	Probably Violative
Other Patterns			
0EM	No Embedded URI-Rs	N/A – intrinsically coherent	Prima Facie Coherent
0NA	No Mementos for URI-R	N/A – resource not archived	Undefined Coherence

[7] Steven Kieffer, Jeremy Avigad, and Harvey Friedman. A language for mathematical language management. Technical Report arXiv:0805.1386v3, 2008.

[8] G. Mohr, M. Stack, I. Rnitovic, D. Avery, and M. Kimpton. Introduction to Heritrix, an archival quality web crawler. In 4th International Web Archiving Workshop, Bath, UK, September 2004.

[9] Kris Carpenter Negulescu. Web archiving @ the Internet Archive. `http://www.digitalpreservation.gov/news/events/ndiipp_meetings/ndiipp10/docs/July21/session09/NDIIPP072110FinalIA.ppt`, 2010.

[10] Michael L. Nelson. Memento-Datetime is not Last-Modified. `http://ws-dl.blogspot.com/2010/11/2010-11-05-memento-datetime-is-not-last.html`, 11 2013.

[11] Marc Spaniol, Dimitar Denev, Arturas Mazeika, Gerhard Weikum, and Pierre Senellart. Data quality in web archiving. In Proceedings of the 3rd Workshop on Information Credibility on the Web, WICOW '09, pages 19–26, 2009.

[12] Marc Spaniol, Arturas Mazeika, Dimitar Denev, and Gerhard Weikum. "Catch me if you can": Visual analysis of coherence defects in web archiving. In The 9th International Web Archiving Workshop (IWAW 2009) Corfu, Greece, September/October, 2009 Workshop Proceedings, pages 27–37, 2009.

[13] Herbert Van de Sompel, Michael Nelson, and Robert Sanderson. HTTP framework for time-based access to resource states—Memento (IETF RFC 7089), December 2013. `http://tools.ietf.org/html/rfc7089`.

[14] Herbert Van de Sompel, Michael L. Nelson, Robert Sanderson, Lyudmila L. Balakireva, Scott Ainsworth, and Harihar Shankar. Memento: Time travel for the Web. Technical Report arXiv:0911.1112, 2009.

[15] Herbert Van de Sompel, Robert Sanderson, Michael Nelson, Lyudmila Balakireva, Harihar Shankar, and Scott Ainsworth. An HTTP-based versioning mechanism for linked data. In Proceedings of Linked Data on the Web Workshop (LDOW2010), April 2010. (Also available as arXiv:1003:3661).

Acknowledgments

This work supported in part by the NSF (IIS 1009392) and the Library of Congress. We are grateful to the Internet Archive for their continued support of Memento access to their archive. Memento is a joint project between the Los Alamos National Laboratory Research Library and Old Dominion University.

I also wish to extend thanks to Yorick Chollet, an intern with Los Alamos National Laboratories, who found cut and paste induced errors in the second version. Predicates 13, 14, 16, 17, 19, and 20 are correct as of version 3.

Appendix A: Last-Modified Is Not Memento-Datetime

Although similar at first look, the *Memento-Datetime* and the *Last-Modified* headers are not equivalent. This appendix briefly describes the difference; we detail the difference in our blog post [10], Although, both are an indicator of when the state of a resource was known, Last-Modified is set by the original server and indicates when the resource was last modified, whereas Memento-Datetime is set by the archive server and indicates when the memento as captured. Consider the URI in the blog post's second case (CD == MD

< LM). Deferencing it yields the headers shown in figure 11. There are four datetime headers, which are indicated in bold.

Compare *Date* to *X-Archive-Original-Date*. The Date header is the datetime the memento was dereferenced—the datetime the archive server responded to the *curl* command. X-Archive-Original-Date, on the other hand, is the datetime that the original resource was dereferenced—the datetime that the original server responded to the archive crawler's request. Assuming that the original resource server and archive crawler have properly-synchronized clocks, X-Archive-Original-Date and Memento-Datetime will be identical or differ by only a second or two. Both Date and Memento-Datetime indicate that the archive crawler captured the memento about 5.5 years ago.

Now look at *X-Archive-Original-Last-Modified*, which is the Last-Modified header provided by the original resource server at the time the crawler captured the memento. At the capture time in March 2008, the original resource server asserted that the page had not changed since December 2002. When original resource servers are functioning properly and Last-Modified is provided, Last-Modified will precede Memento-Datetime, which will precede Date as shown in (21).

$$\text{Last-Modified} < \text{Memento-Datetime} < \text{Date} \quad (21)$$

For original resource servers that set Last-Modified to the current time, Last-Modified will equal Memento-Datetime differ by a second or two, as shown in (22).

$$\text{Last-Modified} \approx \text{Memento-Datetime} < \text{Date} \quad (22)$$

Appendix B: Last-Modified and Dynamic Representations

Although dynamically-generated representations can produce a Last-Modified header, in our experience they typically do not. Also, when a Last-Modified header is produced for a dynamic representation, it is usually set to the current time. Thus, we consider both the absence of Last-Modified or Last-Modified approximately equal to the Date header to be indicative of a dynamic representation. Likewise, the presence of Last-Modified less than Date is indicative of a static representation. Figure 12 shows the headers returned for the CNN home page, which is dynamically generated. Unlike figure 11 from Appendix A, the CNN home page lacks a *X-Archive-Orig-Last-Modified* header.

We expect to empirically test the above assumptions in future work. The end result is expected to be a model using Last-Modified, and possibily other headers, to determine the probability that an archived representation existed at the time the root memento was captured.

Note that the ETag header is defined to have similar, but more flexible, sematics than Last-Modified and can occur independently, they typically co-occur. Thus, ETag generally adds little value in determining static or dynamic representation generation.

```
curl -I 'http://wayback.archive-it.org/927/20080305201649/http://www.nyu.edu/fas/projects/vcb/case_911_FLASHcontent.html'

HTTP/1.1 200 OK
Server: Apache-Coyote/1.1
Memento-Datetime: Wed, 05 Mar 2008 20:16:49 GMT
Link: <http://www.nyu.edu/fas/projects/vcb/case_911_FLASHcontent.html>;
    rel="original",
  <http://wayback.archive-it.org/927/timemap/link/http://www.nyu.edu/fas/projects/vcb/case_911_FLASHcontent.html>;
    rel="timemap"; type="application/link-format",
  <http://wayback.archive-it.org/927/http://www.nyu.edu/fas/projects/vcb/case_911_FLASHcontent.html>;
    rel="timegate",
  <http://wayback.archive-it.org/927/20080305201649/http://www.nyu.edu/fas/projects/vcb/case_911_FLASHcontent.html>;
    rel="first last memento"; datetime="Wed, 05 Mar 2008 20:16:49 GMT"
Set-Cookie: JSESSIONID=8341C31BE000877B3A09DB019FD2824C; Path=/; HttpOnly
X-Archive-Guessed-Charset: UTF-8
X-Archive-Orig-Connection: close
X-Archive-Orig-Content-Length: 4704
X-Archive-Orig-Content-Type: text/html
X-Archive-Orig-ETag: "413f1e-1260-5f9b1700"
X-Archive-Orig-Server: Apache/2.2.3 (Unix) mod\_ssl/2.2.3 OpenSSL/0.9.7d
X-Archive-Orig-Accept-Ranges: bytes
X-Archive-Orig-Last-Modified: Tue, 03 Dec 2002 18:48:28 GMT
X-Archive-Orig-Date: Wed, 05 Mar 2008 20:16:49 GMT
Content-Type: text/html;charset=utf-8
Content-Length: 11483
Date: Thu, 14 Nov 2013 00:34:13 GMT
Connection: close
```

Figure 11: Memento Headers for a Presumably Statically-Generated Representation

```
curl -I 'https://web.archive.org/web/20080317012046/http://www.cnn.com/'

HTTP/1.1 200 OK
Server: Tengine/1.5.1
Date: Fri, 24 Jan 2014 02:25:16 GMT
Content-Type: text/html;charset=utf-8
Content-Length: 116582
Connection: keep-alive
set-cookie: wayback_server=53; Domain=archive.org; Path=/; Expires=Sun, 23-Feb-14 02:25:15 GMT;
Memento-Datetime: Mon, 17 Mar 2008 01:20:46 GMT
Link: <http://www.cnn.com/>; rel="original",
  <http://web.archive.org/web/timemap/link/http://www.cnn.com/>; rel="timemap"; type="application/link-format",
  <http://web.archive.org/web/http://www.cnn.com/>; rel="timegate",
  <http://web.archive.org/web/20000620180259/http://www.cnn.com/>;
    rel="first memento"; datetime="Tue, 20 Jun 2000 18:02:59 GMT",
  <http://web.archive.org/web/20080316150534/http://www.cnn.com/>;
    rel="prev memento"; datetime="Sun, 16 Mar 2008 15:05:34 GMT",
  <http://web.archive.org/web/20080317012046/http://www.cnn.com/>;
    rel="memento"; datetime="Mon, 17 Mar 2008 01:20:46 GMT",
  <http://web.archive.org/web/20080317085057/http://www.cnn.com/>;
    rel="next memento"; datetime="Mon, 17 Mar 2008 08:50:57 GMT",
  <http://web.archive.org/web/20140123231935/http://www.cnn.com/>;
    rel="last memento"; datetime="Thu, 23 Jan 2014 23:19:35 GMT"
X-Archive-Guessed-Charset: cp1252
X-Archive-Orig-vary: Accept-Encoding,User-Agent
X-Archive-Orig-cache-control: max-age=60, private
X-Archive-Orig-content-type: text/html
X-Archive-Orig-x-pad: avoid browser bug
X-Archive-Orig-server: Apache
X-Archive-Orig-content-length: 91164
X-Archive-Orig-expires: Mon, 17 Mar 2008 01:21:44 GMT
X-Archive-Orig-accept-ranges: bytes
X-Archive-Orig-date: Mon, 17 Mar 2008 01:20:46 GMT
X-Archive-Orig-connection: close
X-Archive-Wayback-Perf: [IndexLoad: 571, IndexQueryTotal: 571, RobotsFetchTotal: 5, RobotsRedis: 5,
  RobotsTotal: 5, Total: 777, WArcResource: 71]
Set-Cookie: wb_total_perf=777; Expires=Fri, 24-Jan-2014 02:26:16 GMT; Path=/web/20080317012046/http://www.cnn.com/
X-Archive-Playback: 1
X-Page-Cache: MISS
```

Figure 12: Memento Headers for a Presumably Dynamically-Generated Representation